UNIVERSITÉ DE MONTRÉAL

CENTRE DE RECHERCHE SUR LES TRANSPORTS

CASE POSTALE 6128, SUCCURSALE A, MONTRÉAL, QUÉBEC, H3C 3J7

CUSTOMIZATION VERSUS A GENERAL PURPOSE CODE FOR ROUTING AND SCHEDULING PROBLEMS: A POINT OF VIEW

by

Jean-Marc Rousseau

Cette étude a été publiée grâce à une subvention du fonds F.C.A.R. pour l'aide et le soutien à la recherche.

Université de Montréal
Centre de recherche sur les transports - Publication #556
May 1987

Dépôt légal - 1er trimestre 1988

ABSTRACT

This paper tries to summarize our experience with real life problems of routing and scheduling vehicles. Our experience, coming from a biased sample of these problems, tells us that each problem tends to be different and necessitates a particular attention. We suggest that no single algorithm may be appropriate in all situations and not even in a large number of those and thus that there is a necessity to develop a general approach that can be adapted by experts to each particular situation. We briefly describe problems on which we have worked to exemplify our points.

Key words : routing and scheduling; general comments.

- - - - - - - - - - - -

RÉSUMÉ

Nous tentons dans cette publication de résumer notre expérience sur l'application aux problèmes réels des méthodes d'élaboration de tournées et d'horaires de véhicules. Ces problèmes que nous avons rencontrés à date ont tous tendance à avoir des caractéristiques particulières qui requièrent des méthodes de solution particulières. Nous pensons qu'aucun algorithme général est susceptible d'être applicable de façon satisfaisante à toutes les situations pratiques ni même à un grand nombre de celles-ci. Nous pensons qu'il est nécessaire cependant de développer une approche générale qui peut être adaptée par des experts aux différents contextes particuliers. Nous décrivons brièvement les problèmes sur lesquels nous avons travaillé pour appuyer notre point de vue.

Mots-clés : tournées et horaires; commentaires généraux.

INTRODUCTION

This paper summarizes our experience with real vehicle routing and scheduling problems. Based on our experience with a non-random sample of these problems, we believe that each problem tends to be different and necessitates particular attention. This is however difficult to achieve in the market place because of the lack of knowledge in this area by the various professionals. We suggest that no single algorithm is appropriate in all situations, nor even in a large number of them, and thus that there is a necessity to develop a general approach which can be adapted by experts to each specific situation. We briefly describe problems on which we have worked to exemplify our points.

But first, we tell a story inspired by a real life situation we experienced recently.

A real situation

Once upon a time there was a large distribution company. It owned about 200 stores in Central Canada and it had four depots. It had its own fleet of tractors and trailors but it could also use common carriers when appropriate. It could supply its stores only during operating hours and there were some legal constraints governing the driving time of its drivers. A store's order could be supplied from several depots. Not all depots carried all goods. Split delivery and back orders were allowed, together with multiple trips for the vehicles. It allowed itself up to three working days to supply its stores after receiving an order. Depending on the season, the quantity of goods to be delivered to each store could vary on average from 1/5 to 2/3 of a truckload. The trucks could occasionally pick up goods at manufacturers (backhauling) and eventually deliver goods directly from manufacturers to stores if profitable. Obviously the company wanted to reduce its total operating and delivery costs. It also wanted to minimize the inconvenience to stores of

receiving multiple deliveries. However the exact costs were as usual ill-defined, and linearly approximated.

The company decided recently to give a contract to one of the large consulting firms in computer science, to computerize all its operations (not only the distribution). This large firm, that we will call ABC Inc. is a competent firm in computer science but has no expertise in operations research.

When ABC looked at the distribution problem, they decided to evaluate the packages on the market. Being computer scientists they looked mainly at the specification of the codes available and how easy they were to use. They also obviously looked at their costs and their easiness of implementation.

When they talked to us and explained their problems, we tried to convince them that they needed a customized solution, not because we wanted to make money, but because we thought their problem was too specific for the packages on the market. The algorithms in most packages are geographically based, trying to bring together in the same vehicle clients that are geographically closest to one another. Obviously, vendors of commercial codes insist on the fact that they have an extremely good general purpose algorithm based on state of the art techniques that they cannot reveal for proprietary reasons[1]. However, with a limited number of clients per truck, the full packing of the trucks becomes critical if we want to reduce the number of trucks necessary. In addition, the full use of the multiple trip options should involve more than combining independently generated routes (if possible). Options for generating routes which may be less than full, but that combine well into multiple trip routes should also be considered.

---

[1] We have even heard a vendor use the term "linear programming like" algorithm.

The cost of backorders (delivering part of the load with next delivery) should also be quantified properly together with costs or penalties for split deliveries. The utilization of common carriers for distant deliveries when there is no option for backhauling goods from a manufacturer in the area should be closely looked at. For all these reasons, we believed that no general algorithm could consider all these factors properly and generate the maximum benefit that a customized approach could give. At most, any general package would give the dispatcher the tools to easily modify the original solution proposed. This solution will normally be inadequate or far from optimal, with the added risk that the dispatcher accepts it as optimal because it is produced by a computer.

It is difficult, however, for an operations researcher or expert in routing and scheduling algorithms to convince a computer scientist that even if a specification is there in a package, it is far from obvious that it is dealt with properly and that the underlying algorithm may be totally inadequate for the particularities of their problem. It is also far from obvious that a general purpose package that fully satisfies other well-known corporations will be relevant for a problem with different particularities.

Because we are proposing a customized solution, we obviously do not have it to make a demonstration and prove that a much better solution than the one obtained from a general purpose code can be generated. The cost of even testing a customized solution is greater than the purchase cost of some of the packages on the market. In this context, the temptation for a rational M.I.S. person is to say : "Let us try the general purpose code and then we will see if it is appropriate or not". Only blind faith and confidence in the O.R. people proposing a customized approach can win against this type of argumentation that seems rational at first sight. It is however forgotten that nobody may be able to appreciate how far the

solutions proposed by a general purpose package are from the best obtainable solution. The dispatcher may also get tired of the package and just not use it without telling anybody. In all probability nobody will ever want to take the responsibility of declaring a package recently bought inadequate considering all the costs and efforts invested in trying to make it perform half decently in a very particular context for which it has not been built.[1]

This story has no end as yet. We are still in the midst of it. It could end in several different ways but it is likely that a general purpose code will be purchased. The moral, however, for the O.R. community, is that we are not very successful in transmitting to the engineers, MBA, computer scientists etc. even the most basic notions of O.R. They tend to retain the idea that LP produces optimal solutions, and that probably every algorithm generated by the O.R. community also produces optimal solutions in most situations and it is probably so if the vendor says that his algorithm is optimal or near optimal.

To exemplify this point further, we will describe some of our experience with real problems insisting on the main characteristics of the problem, characteristics which require special attention and customized solutions.

School transportation

In Québec, school children must be transported by the School Board if they live more than one mile away from their school. In rural areas, students must be picked up at their front door and in urban areas they are picked up at stops that are generally located at less than 1/5 of a mile

---

[1] I would like to suggest that this human propensity to spend an order of magnitude more to salvage a bad decision rather than admit the error and start from scratch be called the Mirabel syndrome, in honor of the useless airport in the Montreal area whose contribution to the general economy is negative.

from their home. At first sight, the problem of picking up groups of students at stops is trivial and most packages should be able to do this properly.

But if we look at the problem more deeply, we realize that the problem of routing buses to pick up students with minimal mileage is marginal if considered in relation to the global objective.

The School Boards rent these bus services from private companies. The cost structure is such that the most important cost factor is the number of buses and not the number of miles. The number of buses required is largely a function of the number of school bus routes (picking up students and bringing them to a school) that can be performed by the same bus. This in turn is largely influenced by school starting and ending times. We have shown (see Desrosiers et al., 1986) in a practical context that 30% more buses would be required if all schools start at the same time compared to the situation where schools start at more optimal times relative to each other. In this context what was needed in order of importance was 1) a good algorithm to choose school times, 2) a good algorithm to construct sequences of routing tasks with time windows (there is some flexibility for the arrivals of buses at schools, e.g., 5-20 minutes before school starts), and 3) a good algorithm to construct pick-up routes. Even in this context, the general purpose algorithms used to pick up or deliver goods are not necessarily appropriate because we are transporting people and the shape of the route has special importance. For example, the routes should tend to move in the direction of the schools because parents will complain strongly if the bus passes twice in front of their house or if their children are picked up and then first driven away from schools. In addition, there are constraints for equalizing the load between various buses and reducing the total duration of trips to the schools. At the time we worked on these problems we found the algorithms in

the literature inadequate for this problem and we proposed specially designed algorithms (Chapleau et al., 1984,1985).

This is also a context where we believe that it is necessary to have a detailed network, both to be able to locate the students and evaluate with precision their eligibility for transportation and their walking distance to the bus stop, and to evaluate with precision the travel times of the buses on a network which is relatively sparse and in rural areas full of rivers, mountains and lakes. We have seen situations where the south side of a relatively small mountain is about 15 miles away by road from its north side for a linear distance of about 2 miles. The option of using a full detailed network is not always available in commercial codes.

Postal delivery problems

In the context of postal delivery we have looked at several problems. A first problem is concerned with picking up the contents of mail boxes.

At first sight again this problem is trivial. There is practically no capacity constraint on the vehicles, and because the problem is segmented by period of the day, there is no specific time window except that the length of the routes is limited. However, because there is a relatively large number of mail boxes within a relatively small territory the characteristics of the network become extremely important. In addition, union constraints require that the driver pick up the mail only on the right-hand side of the vehicle (the driver is not allowed to cross the street, for security reasons) and U-turns are generally prohibited. Moreover, many mail boxes are located at street corners and the movement of vehicles is limited after picking up such mail boxes as demonstrated in Figure 1. This implies that the distance from mail box x to mail box y depends also on which mail box z has been serviced before mail box x (see Figure 2).

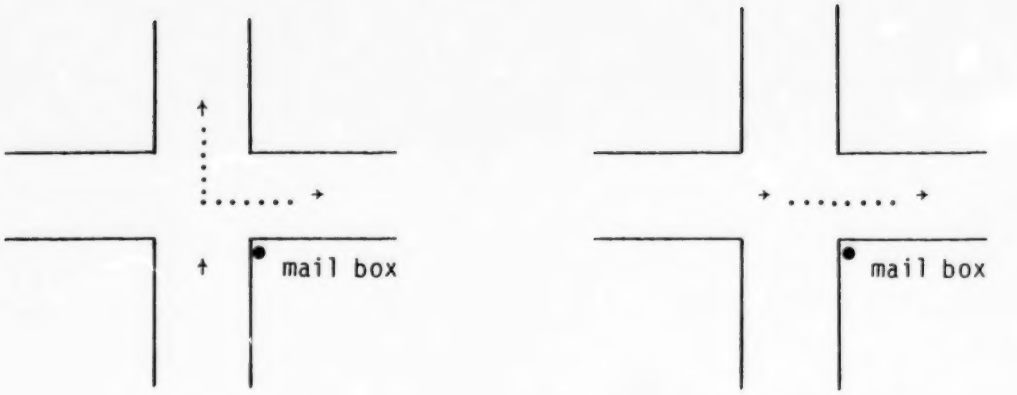

Option 1

The vehicle cannot turn left after picking up at the mail box.

Option 2

The vehicle can only continue straight ahead after picking up at the mail box.

Figure 1 : Turn options when a mail box is located at a street corner.

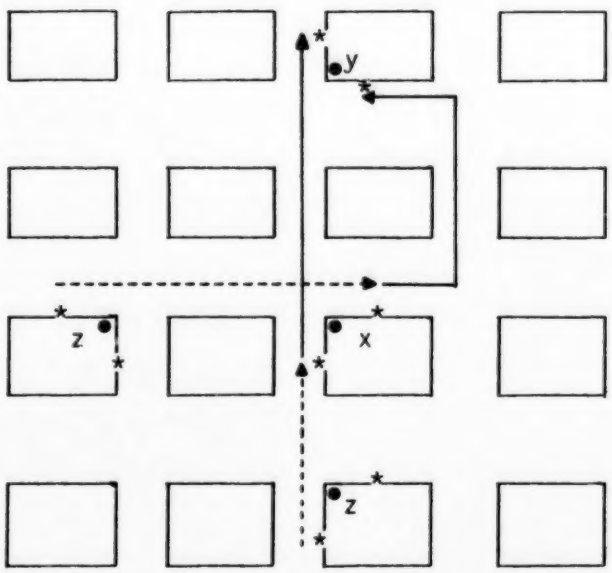

Figure 2 : - actual (•) vs artificial (*) mail boxes;
- the distance (—) between two mail boxes depends on which mail box has been serviced previously (---).

While we are ready to admit that one-way streets and no left turns are really irrelevant for most delivery problems on a large territory, they clearly play a very important role in this context. In fact, if such network characteristics are ignored in the algorithm, the solution produced must generally be totally resequenced manually. The relative error level in the distance between two points is very high if the points are close to each other.

Because the end result must be a set of detailed routes given to truck drivers (including where to turn), the quality of the route proposed must be excellent and hard to improve on manually. Because of this, we also decided to use an exact traveling salesman algorithm to sequence the points on a route using the exact distances calculated on the network. The heuristics used previously were producing nearly optimal solutions but there was some imperfection (crossover) not very important for the total distance but important for the perception of the quality of the routes. In fact we used a version of the traveling salesman algorithm developed by G. Laporte et al. (1985) to take into account mail boxes on street corners (these mail boxes are artificially doubled and located on each of the adjacent arcs; the vehicle is required to service only one of these artificial mail boxes; by doing this the distance between two articifial mail boxes is uniquely determined (see, Figure 2)).

It is worthwhile mentioning that Canada Post Corporation, several years ago, decided to select a package for this problem and did not listen to our comments that a customized system was needed. They bought several copies of a well-known package, and are currently spending a great deal of efforts coding data and trying to make it work. From the previous analysis of the problem we expect that the results will not be very satisfactory.

At Canada Post, we are currently looking at the generation of letter carrier walking routes – a capacitated arc routing problem. For this type of problem there are very few algorithms even published in the literature and, to our knowledge, no package on the market can solve this problem. A paper by Levy and Bodin (1987) describes the work of the authors in this area for the U.S. context. The main characteristic of the problem is that most arcs of the network must be visited and there is very little deadheading in the course of a route (traveling a given arc without delivering mail). All the deadheading occurs at the beginning and at the end of the route and at noon time to go back to the depot for lunch. In the urban areas of Canada, transit buses and taxis are used to transport letter carriers to and from their routes. Balancing of work amongst carriers is also a major goal. Without describing the problem in detail, we can say that because of different ways of transporting the letter carriers to and from the routes, and different union and operating constraints the Canadian and U.S.A. letter-carrier problems may require significantly different algorithms.

Finally we are also studying a routing and scheduling problem for the priority mail system (24-hour delivery system). Pickup points may have fairly restrictive time windows (sometimes ½ hour), but there are no capacity constraints on the vehicles. All vehicles must reach the airport by a specific deadline time. In this context we have tried to adapt a more general algorithm that we have which takes geography into account, but several tests have convinced us that a specific algorithm must also be designed for this context.

Other problems

We have worked on several other problems that we will not describe in detail here. For example, the transportation of the handicapped persons presents one of the most complex routing and scheduling problems because it

includes several capacity constraints (wheel-chair and ambulatory persons), very restrictive time windows on both pickup and drop-off, and obviously the fact that a handicapped person must first be picked up and then dropped off by the same vehicle. This clearly requires a specially-designed algorithm. Other problems include fuel oil delivery, where the main saving does not come from the routing of vehicles but from the careful selection of customers to be serviced during a given day to prevent stockouts and minimize routing costs. Again the problem requires special treatment not available in a general purpose code.

CONCLUSION

In summarizing what we have learned, we would like to emphasize the following points.

Our first point is that in many distribution contexts, the routing aspect and especially the minimization of mileage is not necessarily the most important element of the problem. Organizational problems may generate much greater benefits or must be solved before most benefits can be collected. However, the ability to efficiently produce good routes for vehicles may be essential for the effectiveness of the overall process. For example, in the school bus problem it is important to produce good routes automatically in order that the rest of the scheduling system be applied effectively to achieve all potential savings. Otherwise, the process of elaborating and imputing manually all routes might be too much of a burden and might prevent the use of the rest of the system.

While the simple redesign of letter-carrier routes can bring only few percent savings, several options for modification of the rules involved in making up these routes could generate tremendous savings. On the other hand, any change in organization would necessitate the total redesigning of all letter-carrier walking routes which is an enormous task. While most

savings would come from such a reorganization, the evaluation of impacts and the implementation of any change would need a good system for the design of these routes.

Our second point is that most packages on the market are based on heuristics that give great importance to geography, which can be partially if not totally inadequate for certain types of problems.

Time windows and packing constraints on vehicles, can sometimes almost totally destroy the geographical nature of the problem. In other contexts, the shape of the routes is important. We have mentioned the school bus example. Another example is when the quantity to be delivered is unknown and the vehicles may have to come back to their depot for reloading to continue their routes. In this case, the routes should normally start from the farthest points and converge toward the depot so that if stockouts occur they happen, as much as possible, in the neighborhood of the depot.

We should say however, that the packages found on the market may be useful on a subset of routing and scheduling problems. They should produce good solutions on problems where the detailed network is not important, where the packing aspect is not critical, where time constraints on clients are rather wide with possibly a few exceptions, and no other constraints or operational characteristics are of primal importance for the quality of the solution. At first sight, this may cover a large number of problems, but we are sure that if one gets really into understanding the fundamental nature of each problem there may not be too many practical problems for which they are fully appropriate. Of course, as mentioned previously, the facilities generally included to modify interactively the solution proposed may in all cases be useful to obtain a good solution even if the originally proposed solution is totally inadequate. These features however useful, and even necessary in most cases, are too often an excuse not to invest more in the generation of a high quality initial solution.

Finally, we believe that what is needed is the development of an algorithmic shell that could be instantiated and easily adapted to nearly every context by a routing and scheduling expert. Then a general package could be adapted and have significantly better performance without the necessity of generating a new code each time. Such a shell has been recently proposed by J.-Y. Potvin (Potvin et al., 1987a,b).

Fundamentally, the general heuristic proposed is composed of four basic operators manipulating sets of objects (Figure 3).

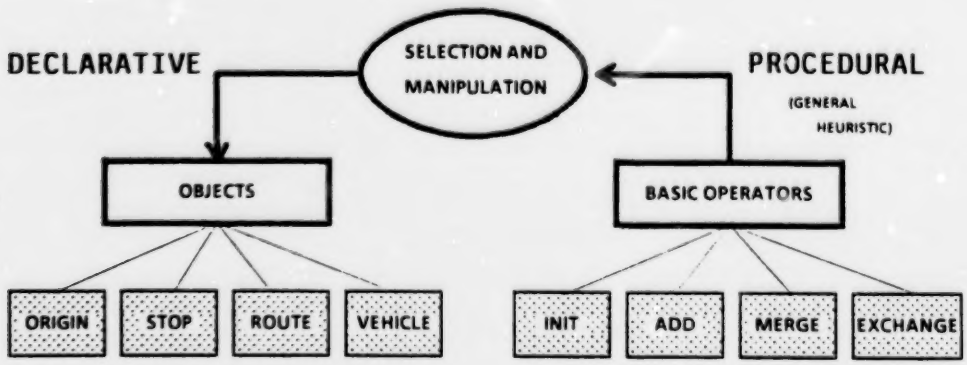

Figure 3 : Objects and basic operators

In the figure, we can identify three major elements:

a) Objects

The objects have attributes providing informations to the operators acting over them. In particular, the major attributes used by the general heuristic are:

i) for the origin : coordinates and number of vehicles available;

ii) for the stops : coordinates, time window and demand;

iii) for the routes : number of stops, length, travel time and total demand;

iv) for the vehicles : speed, capacity and maximal authorized distance or travel time.

b) Basic operators

The objects are manipulated by four basic operators symbolizing distinct phases of the general heuristic. These operators are:

i) INIT : operator for the creation and initialization of new routes;

ii) ADD : operator for the insertion of new stops into existing routes;

iii) MERGE : operator for a two-by-two merging of the routes;

iv) EXCHANGE : operator for the modification of the sequence of stops in the routes.

c) A selection mechanism

Each operator first selects one or more objects and then manipulates these objects. The manipulation process is intrinsically defined in each operator but the selection process is put under user's control by means of selection formulae. These formulae associate a value to the objects involved in the problem's definition. Such values, called OMC for Object Measure Coefficient, allow ordering and selection of objects (for example, we can select the object with minimal or maximal OMC). The most usual selection formulae involve real functions of the following types:

i) $f: E_n \rightarrow R$ (ex: the demand at a stop);

ii) $f: E_r \rightarrow R$ (ex: the length of a route);

iii) $f: E_v \rightarrow R$ (ex: the capacity of some vehicle);

iv) $f: E_n \times E_n \rightarrow R$ (ex: the distance between two nodes);

v) $f: E_n \times E_r \rightarrow R$ (ex: the distance between a node and a route).

where

E_n = the set of nodes in the network;

E_r = the set of routes;

E_v = the set of vehicles.

All operators are independent from each other and can be activated by the user in any desired sequence. The only restriction is obvious: the first operator of the sequence must be INIT (Figure 4). Actually, the user must apply the operators manually but in a near future automatic generation of sequences of properly instantiated operators will be available.

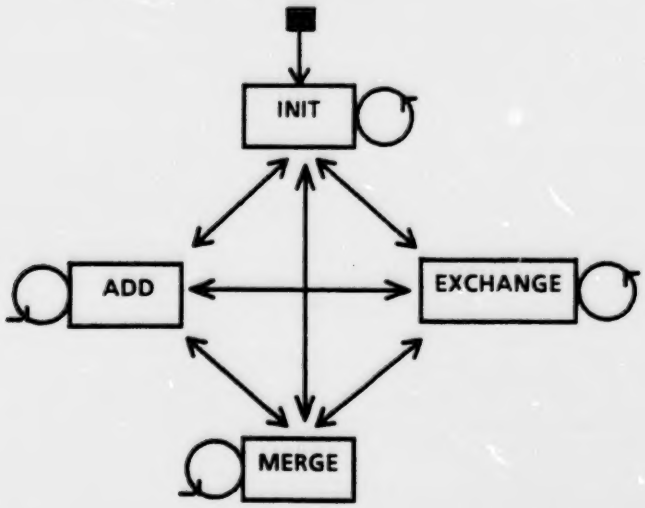

Figure 4 : Activation sequences for the basic operators

In fact, most of the published heuristic algorithms with all their possible variants and mixtures could be easily generated by using appropriate formulae for the selection mechanism and appropriate operators. This general algorithm is at present implemented on a LISP machine and is already helping us in designing new algorithms. We should soon report in the open literature on this recent work. We intend, as part of our development work in this area, to generate the same algorithmic shell on a personal computer so that we can easily design and generate a customized algorithm for each particular situation.

REFERENCES

CHAPLEAU, L., ROUSSEAU, L., FERLAND, J.A. and ROUSSEAU, J.-M. (1985), "Clustering for Routing in Dense Area". European Journal of Operational Research 20:48-57.

CHAPLEAU, L., FERLAND, J.A., LAPALME, G. and ROUSSEAU, J.-M. (1984), "A Parallel Insert Method for the Capacitated Arc Routing Problem". Operations Research Letters 3(2):95-100.

DESROSIERS, J., FERLAND, J.A., ROUSSEAU, J.-M., LAPALME, G. and CHAPLEAU, L. (1986), "TRANSCOL - A Multi-Period School Bus Routing and Scheduling system". In: Delivery of Urban Services, A.J. Swersey and E.J. Ignall (Eds.), Vol. 22, pp. 47-71, in the Series TIMS Studies in the Management Science, Machol, R.E., editor in chief, North-Holland, The Netherlands.

LAPORTE, G., MERCURE, H., and NOBERT, Y. (1985), "Generalized Travelling Salesman Problem Through n Sets of Nodes: The Asymmetrical Case". Publication #399, Centre de recherche sur les transports, Université de Montréal.

LEVY, L., and BODIN (1988), "Scheduling the Postal Carriers for the United States Postal Service: An Application of Arc Partitioning and Routing", in: Vehicle Routing: Methods and Studies, Bruce L. Golden and Arjang A. Assad (Eds), North-Holland Publishing Co., Amsterdam.

POTVIN, J.-Y., LAPALME, G., and ROUSSEAU, J.-M. (1987a), "ALTO: A Computer System for the Design and Experimentation of Routing Algorithms", Publication #525, Centre de recherche sur les transports, Université de Montréal.

POTVIN, J.-Y. (1987b), "Un système informatique pour le développement et l'expérimentation d'algorithmes de génération de tournées". Thèse de Ph.D., département d'informatique et de recherche opérationnelle, publication #522, Centre de recherche sur les transports, Université de Montréal.